

# Hybrid Conditional Gradient - Smoothing Algorithms with Applications to Sparse and Low Rank Regularization

A. Argyriou

École Centrale Paris, Center for Visual Computing  
andreas.argyriou@ecp.fr

M. Signoretto

KU Leuven, ESAT-STADIUS  
Kasteelpark Arenberg 10, B-3001 Leuven, Belgium  
marco.signoretto@esat.kuleuven.be

J. Suykens

KU Leuven, ESAT-STADIUS  
Kasteelpark Arenberg 10, B-3001 Leuven, Belgium  
johan.suykens@esat.kuleuven.be

## 1 Introduction

*Conditional gradient* methods are old and well studied optimization algorithms. Their origin dates at least to the 50's and the Frank-Wolfe algorithm for quadratic programming [18] but they apply to much more general optimization problems. General formulations of conditional gradient algorithms have been studied in the past and various convergence properties of these algorithms have been proven. Moreover, such algorithms have found application in many fields, such as optimal control, statistics, signal processing, computational geometry and machine learning. Currently, interest in conditional gradient methods is undergoing a revival because of their computational advantages when applied to certain large scale optimization problems. Such examples are *regularization* problems involving sparsity or low rank constraints, which appear in many widely used methods in machine learning.

Inspired by such algorithms, in this chapter we study a first-order method for solving certain convex optimization problems. We focus on problems of the form

$$\min \{f(x) + g(Ax) + \omega(x) : x \in \mathcal{H}\} .$$

over a real Hilbert space  $\mathcal{H}$ . We assume that  $f$  is a convex function with *Hölder continuous gradient*,  $g$  a *Lipschitz continuous* convex function,  $A$  a bounded linear operator and  $\omega$  a convex function defined over a *bounded domain*. We also assume that the computational operations available are the *gradient* of  $f$ , the *proximity operator* of  $g$  and a *subgradient of the convex conjugate*  $\omega^*$ .<sup>1</sup> A particularly common type of problems covered by (1.1) is

$$\min \{f(x) + g(Ax) : x \in \mathcal{C}\} ,$$

where  $\mathcal{C}$  is a bounded, closed, convex subset of  $\mathcal{H}$ . Common such examples are regularization problems with one or more penalties in the objective (as the term  $g \circ A$ ) and one penalty as a constraint described by  $\mathcal{C}$ .

Before presenting the algorithm, we review in Section 3 a generic conditional gradient algorithm which has been well studied in the past. This standard algorithm can be used for solving problems of the form (1.1) whenever  $g = 0$ . However, the conditional gradient algorithm cannot handle problems with a nonzero term  $g$ , because it would require computation of a subgradient of a composite convex conjugate function, namely a subgradient of  $(g \circ A + \omega)^*$ . In many cases of interest, there is no simple rule for such subgradients and the computation itself requires an iterative algorithm.

Thus, in Section 4 we discuss an alternative approach that combines ideas from both conditional gradient algorithms and smoothing proximal algorithms, such as Nesterov smoothing. We call the resulting algorithm a *hybrid conditional gradient - smoothing* algorithm, in short HCGS. This approach involves smoothing the  $g$  term, that is, approximating  $g$  with a function whose gradient is Lipschitz continuous. Besides this modification, HCGS is similar to the conditional gradient algorithm. We show that, for suitable choices of the smoothing parameter, the estimates of the objective in HCGS converge to the minimum of (1.1). Moreover, the convergence rate is of the order of  $\mathcal{O}(\frac{1}{\varepsilon^2})$  iterations for attaining an accuracy of  $\varepsilon$  in terms of the objective. We do not claim originality, however, since similar theoretical results have appeared in recent work by Lan [34].

Our main focus is on highlighting applications of the hybrid approach on certain applications of interest. To demonstrate the applicability of HCGS to regularization problems from machine learning, we present simulations on matrix problems with *simultaneous sparsity and low rank* penalizations. Examples of such applications are *graph denoising*, *link prediction* in social networks, *covariance estimation* and *sparse PCA*. Each of these problems involves two penalties, an elementwise  $\ell_1$  norm to promote sparsity, and a *trace norm* to promote low rank. Standard algorithms may not be practical in high dimensional problems of this type. As mentioned above, standard conditional gradient methods require a subgradient computation of a complicated function, whereas proximal algorithms or subgradient based algorithms require an expensive singular value decomposition per iteration. In contrast, HCGS requires only computation of dominant singular vectors, which is more practical by means of the

---

<sup>1</sup>For the precise assumptions required, see Assumptions 4.1, 4.3.

power iteration. Thus, even though HCGS exhibits a slower asymptotic rate of convergence than conditional gradient algorithms, Nesterov's method or the forward-backward algorithm, it scales much better to large matrices than these methods.

## 2 Preliminaries from Convex Analysis

Throughout the chapter,  $\mathcal{H}$  is a real Hilbert space endowed with norm  $\|\cdot\|$  and inner product  $\langle \cdot, \cdot \rangle$ . As is standard in convex analysis, we consider extended value functions  $f : \mathcal{H} \rightarrow (-\infty, +\infty]$  which can take the value  $+\infty$ . With this notation, constraints can be written as indicator functions that take the zero value inside the feasible set and  $+\infty$  outside.

**Definition 2.1.** The domain of a function  $f : \mathcal{H} \rightarrow (-\infty, +\infty]$  is defined as the set  $\text{dom } f = \{x \in \mathcal{H} : f(x) < +\infty\}$ .

**Definition 2.2.** The function  $f : \mathcal{H} \rightarrow (-\infty, +\infty]$  is called *proper* if  $\text{dom } f \neq \emptyset$ .

**Definition 2.3.** The set of proper lower semicontinuous convex functions from  $\mathcal{H}$  to  $(-\infty, +\infty]$  is denoted by  $\Gamma_0(\mathcal{H})$ .

**Definition 2.4.** Let  $f : \mathcal{H} \rightarrow [-\infty, +\infty]$ . The *convex conjugate* of  $f$  is the function  $f^* : \mathcal{H} \rightarrow [-\infty, +\infty]$  defined as

$$f^*(x) = \sup\{\langle u, x \rangle - f(u) : u \in \mathcal{H}\}$$

for every  $x \in \mathcal{H}$ .

**Theorem 2.1.** (Fenchel-Moreau) [4, Thm. 13.32]. Every function  $f \in \Gamma_0(\mathcal{H})$  is biconjugate,

$$f^{**} = f.$$

Moreover,  $f^* \in \Gamma_0(\mathcal{H})$ .

**Definition 2.5.** Let  $f : \mathcal{H} \rightarrow (-\infty, +\infty]$  be proper. A *subgradient* of  $f$  at  $x \in \mathcal{H}$  is a vector  $u \in \mathcal{H}$  satisfying

$$\langle u, y - x \rangle + f(x) \leq f(y)$$

for every  $y \in \mathcal{H}$ . The set of subgradients of  $f$  at  $x$  is called the *subdifferential* of  $f$  at  $x$  and is denoted as  $\partial f(x)$ .

**Proposition 2.2.** Let  $f : \mathcal{H} \rightarrow (-\infty, +\infty]$  be proper and  $x \in \mathcal{H}$ . Then  $\partial f(x) \neq \emptyset$  implies  $x \in \text{dom } f$ .

**Theorem 2.3.** [4, Thm. 16.23]. Let  $f \in \Gamma_0(\mathcal{H})$ ,  $x \in \mathcal{H}$ ,  $u \in \mathcal{H}$ . Then

$$\begin{aligned} x \in \partial f(x) &\iff x \in \partial f^*(u) \iff f(x) + f^*(u) = \langle x, u \rangle, \\ (\partial f)^{-1} &= \partial f^*. \end{aligned}$$

**Theorem 2.4.** [4, Thm. 16.37]. Let  $f \in \Gamma_0(\mathcal{H})$ ,  $\mathcal{K}$  a Hilbert space,  $g \in \Gamma_0(\mathcal{K})$  and  $A : \mathcal{H} \rightarrow \mathcal{K}$  a bounded linear operator. If  $\text{dom } g = \mathcal{K}$  then

$$\partial(f + g \circ A) = \partial f + A^* \circ \partial g \circ A .$$

**Theorem 2.5.** [4, Thm. 16.2]. Let  $f : \mathcal{H} \rightarrow (-\infty, +\infty]$  be proper. Then

$$\text{argmin } f = \{x \in \mathcal{H} : 0 \in \partial f(x)\} .$$

**Definition 2.6.** The function  $f \in \Gamma_0(\mathcal{H})$  is called  $(p, L)$ -smooth, where  $L > 0$ ,  $p \in (0, 1]$ , if  $f$  is Fréchet differentiable on  $\mathcal{H}$  with a Hölder continuous gradient,

$$\|\nabla f(x) - \nabla f(y)\| \leq L \|x - y\|^p \quad \forall x, y \in \mathcal{H} .$$

The case  $p = 1$  corresponds to functions with Lipschitz continuous gradient, which appear frequently in optimization. The following lemma is sometimes called the *descent lemma* [4, Cor. 18.14].

**Lemma 2.6.** If the function  $f \in \Gamma_0(\mathcal{H})$  is  $(p, L)$ -smooth then

$$f(x) \leq f(y) + \langle x - y, \nabla f(y) \rangle + \frac{L}{p+1} \|x - y\|^{p+1} \quad \forall x, y \in \mathcal{H} . \quad (2.1)$$

**Theorem 2.7.** Let  $\omega \in \Gamma_0(\mathcal{H})$ . Then  $\text{dom } \omega$  is contained in the ball of radius  $\rho \in \mathbb{R}_+$ , if and only if the convex conjugate  $\omega^*$  is  $\rho$ -Lipschitz continuous on  $\mathcal{H}$ .

**Proof.** See [47], or [48, Cor. 13.3.3] for a finite-dimensional version.  $\square$

**Corollary 2.8.** If the function  $\omega \in \Gamma_0(\mathcal{H})$  has bounded domain then  $\partial\omega^*$  is nonempty everywhere on  $\mathcal{H}$ .

**Proof.** Follows from Theorem 2.7 and [4, Prop. 16.17].  $\square$

### 3 Generalized Conditional Gradient Algorithm

In this section, we review briefly the conditional gradient algorithm in one of its many formulations. We focus on convex optimization problems of a general type and discuss how a generalized conditional gradient algorithm applies to such problems. We should note that this algorithm is not the most generic formulation that has been studied – see, for example, [34, 13] – but it covers a broad variety of optimization problems in machine learning.

Specifically, we consider the optimization problem

$$\min \{f(x) + \omega(x) : x \in \mathcal{H}\} \quad (3.1)$$

where we make the following assumptions.

---

**Algorithm 1** Generalized conditional gradient algorithm.

---

**Input**  $x_1 \in \text{dom } \omega$   
**for**  $k = 1, 2, \dots$  **do**  
 $y_k \leftarrow \text{an element of } \partial\omega^*(-\nabla f(x_k))$  (I)  
 $x_{k+1} \leftarrow (1 - \alpha_k)x_k + \alpha_k y_k$  (II)  
**end for**

---

**Assumption 3.1.**

- $f, \omega \in \Gamma_0(\mathcal{H})$
- $f$  is  $(1, L)$ -smooth
- $\text{dom } \omega$  is bounded, that is, there exists  $\rho \in \mathbb{R}_{++}$  such that  $\|x\| \leq \rho, \forall x \in \text{dom } \omega$ .

**Remark 3.2.** Under Assumption 3.1, the problem (3.1) admits a minimizer. The reason is that, since  $\text{dom } \omega$  is bounded,  $\lim_{\|x\| \rightarrow +\infty} \frac{f(x) + \omega(x)}{\|x\|} = +\infty$  (supercoercivity).

The Fenchel dual problem associated with (3.1) is

$$\max \{-f^*(-z) - \omega^*(z) : z \in \mathcal{H}\} . \quad (3.2)$$

Due to Fenchel's duality theorem and the fact that  $\text{dom } f = \mathcal{H}$ , the duality gap equals zero and the maximum in (3.2) is attained [4, Thm. 15.23].

Algorithm 1 has been used frequently in the past for solving problems of the type (3.1). It is a generalization of algorithms such as the Frank-Wolfe algorithm for quadratic programming [18, 21] and conditional gradient algorithms [32, 14, 15]. Algorithm 1 applies to the general setting of convex optimization problems of the form (3.1) which satisfy Assumption 3.1. In such general forms, the algorithm has been known and studied for a long time in control theory and several of its convergence properties have been obtained [31, 32, 15, 14, 16]. More recently, interest in the family of conditional gradient algorithms has been revived, especially in theoretical computer science, machine learning, computational geometry and elsewhere [24, 23, 29, 3, 20, 54, 56, 27, 10, 22, 33, 19]. Some of these algorithms have appeared independently in various fields, such as statistics and signal processing, under different names and various guises. For example, it has been observed that conditional gradient methods are related to boosting, greedy methods for sparse problems [10, 51] and to orthogonal matching pursuit [28, 27]. Some very recent papers [3, 53] show an equivalence to the optimization method of mirror descent, which we discuss briefly in Section 3.2.

One reason for the popularity and the revival of interest in conditional gradient methods has been their applicability to large scale problems. This advantage is evident, for example, in comparison to proximal methods – see [11] and references therein – and especially in optimization problems involving matrices. Conditional gradient methods generally trade off a slower convergence rate

(number of iterations) for lower complexity of each iteration step. The accelerated proximal gradient methods [43] benefit from the “optimal”  $\mathcal{O}\left(\sqrt{\frac{1}{\varepsilon}}\right)$  rate (where  $\varepsilon$  is the accuracy with respect to the optimization objective), whereas conditional gradient methods exhibit a slower  $\mathcal{O}\left(\frac{1}{\varepsilon}\right)$  rate. On the other side, each step in the proximal methods requires computation of the *proximity operator* [37, 4] (see Section 4), which in some cases can be particularly costly. For example, the proximity operator of the *trace norm* of a matrix  $X \in \mathbb{R}^{d \times n}$ ,

$$\|X\|_{tr} = \sum_{i=1}^{\min\{d,n\}} \sigma_i(X) ,$$

where  $\sigma_i(X)$  denote the singular values of  $X$ , requires computation of a complete singular value decomposition. In contrast, a conditional gradient method need only compute a dominant pair of left-right singular vectors, and such a computation scales better to large matrices [29].

In general, as Algorithm 1 indicates, conditional gradient methods require computation of *dual subgradients*. Often, this is a much less expensive operation than projection or the proximity operation. In other cases, proximity operations may not be feasible in a finite number of steps, whereas dual subgradients are easy to compute. An obvious such case is  $\ell_p$  or Schatten- $\ell_p$  regularization – see [28]. Other cases of interest occur when  $\omega$  is the conjugate of a *max-function*. Then the dual subgradient could be fast to compute while the proximity operation may be complex.

Finally, another advantage of conditional gradient methods is that they build their estimate of the solution incrementally. This implies that, in earlier iterations, time and space costs will be low and that the algorithm may be stopped once an estimate of the desired parsimony is obtained (this could be, for example, a vector of certain sparsity or a matrix of certain rank). Proximal methods, in contrast, do not necessarily obtain the desired parsimony until later iterations (and even then it is not “exact”).

Formulation (3.1) covers many optimization problems studied so far in the conditional gradients literature and provides a concise description of variational problems amenable to the standard conditional gradient algorithm. In Section 4 we extend the applicability to problems with multiple penalties, by combining conditional gradients and smoothing techniques.

We remark that formulation (3.1) is valid in a generalized Hilbert space setting, so that it can be applied to infinite dimensional problems. This is particularly useful for *kernel methods* in machine learning, for example, kernelized support vector machines or structured SVMs [52] and nuclear or Schatten- $\ell_p$  regularization of operators [1].

To motivate Algorithm 1, consider the convex optimization problem (3.1). By Theorems 2.4 and 2.5,  $\hat{x} \in \mathcal{H}$  is a minimizer of (3.1) if and only if

$$0 \in \nabla f(\hat{x}) + \partial\omega(\hat{x})$$

or, equivalently,

$$-\nabla f(\hat{x}) \in \partial\omega(\hat{x}) \iff \hat{x} \in \partial\omega^*(-\nabla f(\hat{x})), \quad (3.3)$$

where we have used Theorem 2.3. Thus, step (I) in Algorithm 1 reflects the fixed point equation (3.3). However,  $\partial\omega^*(-\nabla f(x_k))$  is not a singleton in general and some elements of this set may be far from the minimizers of the problem. Hence step (II), which weighs the new estimate with past ones, is necessary. With any affine weighting like that of step (II), the fixed point equation (3.3) still holds.

**Remark 3.3.** Algorithm 1 is *well defined*, since the subdifferential at step (I) is always nonempty, due to Corollary 2.8 and Assumption 3.1.

Finally, let us note that several variants of Algorithm 1 are possible, in the spirit of the extensive literature on conditional gradient methods. For example, there are various techniques (like line search) for the choice of coefficients  $\alpha_k$ , more of the past iterates may be used in (II) and so on.

### 3.1 Convergence Rate

**Theorem 3.1.** *If, for every  $k \in \mathbb{N}$ ,  $\alpha_k \in [0, 1]$ , then  $x_k \in \text{dom } \omega$  and*

$$f(x_{k+1}) + \omega(x_{k+1}) - f(x) - \omega(x) \leq (1 - \alpha_k)(f(x_k) + \omega(x_k) - f(x) - \omega(x)) + 2\alpha_k^2 L\rho^2$$

for every  $x \in \text{dom } \omega$ ,  $k \in \mathbb{N}$ .

Theorem 3.1 implies an  $\mathcal{O}(\frac{1}{k})$  convergence rate with respect to the objective values  $f(x_k) + \omega(x_k) - f(\hat{x}) - \omega(\hat{x})$ , where  $\hat{x}$  is a minimizer of (3.1). This rate can be attained, for example, with the choice  $\alpha_k = \frac{2}{k+1}$ .

**Corollary 3.2.** *If  $\alpha_k = \frac{2}{k+1}$ , for every  $k \in \mathbb{N}$ , then*

$$f(x_{k+1}) + \omega(x_{k+1}) - f(x) - \omega(x) \leq \frac{8L\rho^2}{k+1} \quad (3.4)$$

for every  $x \in \text{dom } \omega$ ,  $k \in \mathbb{N}$ .

See [14, 15, 29, 10, 3] and references therein for these and related results, as well as for bounds involving the duality gap estimates. It is also known that the lower bound for conditional gradient and similar algorithms is of the same order [9, 28, 34].

### 3.2 Connections to Mirror Descent and Gradient Descent

It has been observed recently [3, 53] that the conditional gradient algorithm is equivalent to a *mirror descent* algorithm in the dual. The basic mirror descent algorithm [5, 38, 30] may be written as the iteration

$$x_{k+1} \leftarrow \text{an element of } x_k - t_k \partial\varphi(\nabla\psi^*(x_k)), \quad (3.5)$$

where  $t_k > 0$  are step sizes,  $\psi$  is strongly convex on a closed convex set  $C$  and  $\varphi$  is convex and Lipschitz continuous on  $C$ . Setting  $\omega = (\varphi \circ (-I))^*$ , where  $I$  denotes the identity operator, and  $f = \psi^*$ , algorithm (3.5) rewrites as a variant of Algorithm 1 (in which the update is not a convex combination). The set  $C$  can be viewed as the domain of  $\omega$ .

Consequently, when (3.1) is a proximity computation (that is, when  $f = \frac{1}{2\beta} \|\cdot\|^2$ ,  $\beta > 0$ ) the conditional gradient algorithm 1 is equivalent to a *subgradient descent* in the dual. In such cases  $\nabla f = \frac{1}{\beta} I$  and Algorithm 1 becomes

$$x_{k+1} \in (1 - \alpha_k)x_k + \alpha_k \partial \omega^* \left( -\frac{1}{\beta} x_k \right) .$$

Letting  $h = \left( \omega^* \circ \left( -\frac{1}{\beta} I \right) \right)^*$ , by the chain rule (Theorem 2.4) this iteration is equivalent to

$$x_{k+1} \in x_k - \alpha_k (I + \beta \partial h^*)(x_k) .$$

In particular, when  $\omega$  is  $\mu$ -strongly convex (and hence  $\omega^*$  is  $(1, \frac{1}{\mu})$ -smooth) and  $\alpha_k \leq \frac{\mu\beta}{1+\mu\beta}$  for every  $k \in \mathbb{N}$ , the above iteration is equivalent to a *proximal point* algorithm [11, 17, 49] because of [4, Thm. 18.15]. Note that not all cases of subgradient descent are covered, since  $\omega$  should have bounded domain, implying that the dual objective function should be a quadratic perturbation of a Lipschitz continuous function.

## 4 Hybrid Conditional Gradient - Smoothing Algorithm

We now introduce Algorithm 2, an extension of conditional gradient methods to optimization problems on bounded domains which contain smooth and Lipschitz continuous terms.

### 4.1 Description of the Hybrid Algorithm

Formally, we consider the class of optimization problems of the form

$$\min \{f(x) + g(Ax) + \omega(x) : x \in \mathcal{H}\} \quad (4.1)$$

where we make the following assumptions:

**Assumption 4.1.**

- $f, \omega \in \Gamma_0(\mathcal{H})$
- $g \in \Gamma_0(\mathcal{K})$ ,  $\mathcal{K}$  is a Hilbert space
- $A : \mathcal{H} \rightarrow \mathcal{K}$  is a bounded linear operator
- $f$  is  $(p, L_f)$ -smooth



- $g$  is  $L_g$ -Lipschitz continuous on  $\mathcal{K}$
- $\text{dom } \omega$  is bounded, that is, there exists  $\rho \in \mathbb{R}_{++}$  such that  $\|x\| \leq \rho, \forall x \in \text{dom } \omega$

**Remark 4.2.** Under Assumption 4.1, problem (4.1) admits a minimizer. As in Remark 3.2, the reason is growth of the objective function at infinity (the objective equals  $+\infty$  outside the feasible set, which is bounded).

In order for the algorithm to be practical, we require that

**Assumption 4.3.**

- the gradient of  $f$  is simple to compute at every  $x \in \mathcal{H}$ ,
- a subgradient of  $\omega^*$  is simple to compute at every  $x \in \mathcal{H}$ ,
- the proximity operator of  $\beta g$  is simple to compute for every  $\beta > 0, x \in \mathcal{H}$ .

The *proximity operator* was introduced by Moreau [37] as the (unique) minimizer

$$\text{prox}_g(x) = \operatorname{argmin} \left\{ \frac{1}{2} \|x - u\|^2 + g(u) : u \in \mathcal{H} \right\}.$$

For a review of the numerous applications of proximity operators to optimization, see, for example, [12, 11] and references therein.

The following are some examples of optimization problems that belong to the general class (4.1).

*Example 4.1.* Regularization with two norm penalties:

$$\min \{ f(x) + \lambda \|x\|_a : \|x\|_b \leq B, x \in \mathbb{R}^d \}$$

where  $f$  is  $(p, L_f)$ -smooth,  $\lambda > 0$  and  $\|\cdot\|_a, \|\cdot\|_b$  can be any norms on  $\mathbb{R}^d$ .

*Example 4.2.* Regularization with a linear composite penalty and a norm:

$$\min \{ f(x) + \lambda \|Ax\|_a : \|x\|_b \leq B, x \in \mathbb{R}^d \}$$

where  $f$  is  $(p, L_f)$ -smooth,  $\lambda > 0$ ,  $\|\cdot\|_a, \|\cdot\|_b$  are norms on  $\mathbb{R}^\delta, \mathbb{R}^d$ , respectively, and  $A \in \mathbb{R}^{\delta \times d}$ .

*Example 4.3.* Regularization with multiple linear composite penalties and a norm:

$$\min \left\{ f(x) + \sum_{i=1}^n \lambda_i \|A_i x\|_{a_i} : \|x\|_b \leq B, x \in \mathbb{R}^d \right\}$$

where  $f$  is  $(p, L_f)$ -smooth and, for all  $i \in \{1, \dots, n\}$ ,  $\lambda_i > 0$ ,  $\|\cdot\|_{a_i}, \|\cdot\|_b$  are norms on  $\mathbb{R}^{\delta_i}, \mathbb{R}^d$ , respectively, and  $A_i \in \mathbb{R}^{\delta_i \times d}$ . Such problems can be seen as special cases of Example 4.2 by applying the classical direct sum technique,

$$\delta = \sum_{i=1}^n \delta_i, A = \begin{pmatrix} A_1 \\ \vdots \\ A_n \end{pmatrix}, \|(v_i)_{i=1}^n\|_a = \sum_{i=1}^n \lambda_i \|v_i\|_{a_i}.$$

---

**Algorithm 2** Hybrid conditional gradient - smoothing algorithm.

---

**Input**  $x_1 \in \text{dom } \omega$   
**for**  $k = 1, 2, \dots$  **do**  
 $z_k \leftarrow -\nabla f(x_k) - \frac{1}{\beta_k} A^* A x_k + \frac{1}{\beta_k} A^* \text{prox}_{\beta_k g}(A x_k)$  (I)  
 $y_k \leftarrow \text{an element of } \partial \omega^*(z_k)$  (II)  
 $x_{k+1} \leftarrow (1 - \alpha_k) x_k + \alpha_k y_k$  (III)  
**end for**

---

We propose to solve problems like the above with Algorithm 2.<sup>2</sup> We call it a *hybrid conditional gradient - smoothing* algorithm (HCGS in short), because it involves a smoothing of function  $g$  with parameter  $\beta_k$ . For any  $\beta > 0$ , the  $\beta$ -smoothing of  $g$  is the Moreau envelope of  $g$ , that is, the function  $g_\beta$  defined as

$$g_\beta(x) := \min \left\{ \frac{1}{2\beta} \|x - u\|^2 + g(u) : u \in \mathcal{H} \right\} \quad \forall x \in \mathcal{H}. \quad (4.2)$$

The function  $g_\beta$  is a smooth approximation to  $g$  (in fact, the best possible approximation of  $\frac{1}{\beta}$  smoothness), as summarized in the following lemmas from the literature.

**Lemma 4.1** (Proposition 12.29 in [4]). *Let  $g \in \Gamma_0(\mathcal{H})$ ,  $\beta > 0$ . Then  $g_\beta$  is  $(1, \frac{1}{\beta})$ -smooth and its gradient can be obtained from the proximity operator of  $g$  as:  $\nabla g_\beta(x) = \frac{1}{\beta} (x - \text{prox}_{\beta g}(x))$ .*

**Lemma 4.2.** *Let  $g \in \Gamma_0(\mathcal{H})$  be  $L_g$ -Lipschitz continuous and  $\beta > 0$ . Then*

- $g_\beta \leq g \leq g_\beta + \frac{1}{2}\beta L_g^2$
- if  $\beta \geq \beta' > 0$ , then  $g_\beta \leq g_{\beta'} \leq g_\beta + \frac{1}{2}(\beta - \beta') L_g^2$ .

**Proof.** Define  $\Psi_x(u) = \frac{1}{2\beta} \|x - u\|^2 + g(u)$ . We have  $g_\beta(x) = \min_u \Psi_x(u) \leq \Psi_x(x) = g(x)$ , and this proves the left hand side of the first property. For the other side of the inequality, we have

$$\Psi_x(u) = \frac{1}{2\beta} \|x - u\|^2 + g(u) - g(x) + g(x) \geq \frac{1}{2\beta} \|x - u\|^2 + g(x) - L_g \|x - u\|$$

where we have used the Lipschitz property of  $g$ . This implies that

$$g_\beta(x) = \min_u \Psi_x(u) \geq g(x) + \inf_u \left( \frac{1}{2\beta} \|x - u\|^2 - L_g \|x - u\| \right) = g(x) - \frac{1}{2}\beta L_g^2.$$

The second property follows from the first one and  $(g_{\beta'})_{\beta - \beta'} = g_\beta$  (Proposition 12.22 in [4]).  $\square$

---

<sup>2</sup> Algorithm 2 is well-defined (see Remark 3.3).

Thus, the smoothing parameter  $\beta$  controls the tradeoff between the smoothness and the quality of approximation.

At each iteration, the hybrid Algorithm 2 computes the gradient of the *smoothed* part  $f + g_{\beta_k} \circ A$ , where  $\beta_k$  is the adaptive smoothing parameter. By Lemma 4.1 and the chain rule, its gradient equals

$$\nabla(f + g_{\beta_k} \circ A)(x) = \nabla f(x) + \frac{1}{\beta_k} A^* (Ax - \text{prox}_{\beta_k g}(Ax)) .$$

The function  $f$  is  $(p, L_f)$ -smooth and the function  $g_{\beta_k} \circ A$  is  $(1, \frac{1}{\beta_k} \|A\|^2)$ -smooth. By selecting  $\beta_k$  that approaches 0 as  $k$  increases, we ensure that  $g_{\beta_k}$  approaches  $g$ .

Algorithm 2 can be viewed as an extension of conditional gradient algorithms and of Algorithm 1. But besides conditional gradient methods, the algorithm also exploits ideas from proximal algorithms obtained by smoothing Lipschitz terms in the objective. These methods are primarily due to Nesterov and have been successfully applied to many problems [42, 41, 44, 25, 7]. The smoothing we apply here is a type of Moreau envelope as in the variational problem (4.2), which is connected to Nesterov smoothing – see, for example, [44, 7].

However, unlike Nesterov’s smoothing and other proximal methods, in our method we choose not to smooth function  $\omega$ , or apply any other proximity-like operation to it. We do this because computation of the proximity operator of  $\omega$  is not available in the settings which we consider here.<sup>3</sup> For example, if  $\omega$  expresses a trace norm constraint, the proximity computation requires a *full* singular value decomposition, which does not scale well with the size of the matrix. In contrast, the dual subgradient requires only computation of a single pair of dominant singular vectors and this is feasible even for very large matrices using the power method or Lanczos algorithms.

## 4.2 Convergence Rate

A bound on the convergence rate of the objective function can be obtained by first bounding the convergence of the smoothed objective in a recursive way. The required number of iterations is a function of  $p$  and  $\varepsilon$ , where  $p$  is the smoothing exponent of  $f$  and  $\varepsilon$  is the accuracy in terms of the objective function. Regarding the proof technique, we should note that the HCGS Algorithm 2 and the proof of its convergence properties are mostly related to conditional gradient methods. On the other side, the proof technique does not share similarities with proximal methods such as ISTA or FISTA [6, 39, 40].

**Theorem 4.3.** *Suppose that, for every  $k \in \mathbb{N}$ ,  $\alpha_k \in [0, 1]$  and  $\beta_k \geq \beta_{k+1} > 0$ .*

---

<sup>3</sup>Note that Nesterov’s smoothing would require  $\omega$  to be Lipschitz continuous, and hence it does not apply directly to the case of bounded  $\text{dom } \omega$  but to a similar regularization problem with  $\omega$  as a penalty in the objective.

Let  $F = f + g \circ A$ ,  $F_k = f + g_{\beta_k} \circ A$ . Then  $x_k \in \text{dom } \omega$ , for every  $k \in \mathbb{N}$ , and

$$\begin{aligned} F_{k+1}(x_{k+1}) + \omega(x_{k+1}) - F(x) - \omega(x) &\leq (1 - \alpha_k)(F_k(x_k) + \omega(x_k) - F(x) - \omega(x)) \\ &\quad + \frac{(2\rho)^{p+1}L_f}{p+1}\alpha_k^{p+1} + 2\|A\|^2\rho^2\frac{\alpha_k^2}{\beta_k} + \frac{1}{2}(\beta_k - \beta_{k+1})L_g^2 \end{aligned} \quad (4.3)$$

for every  $x \in \text{dom } \omega$ ,  $k \in \mathbb{N}$ .

**Proof.** For every  $k \in \mathbb{N}$ , we apply the descent lemma 2.6 twice to obtain that

$$\begin{aligned} f(x_{k+1}) &\leq f(x_k) + \langle \nabla f(x_k), x_{k+1} - x_k \rangle + \frac{L_f}{p+1}\|x_{k+1} - x_k\|^{p+1} \\ &= f(x_k) + \alpha_k \langle \nabla f(x_k), y_k - x_k \rangle + \frac{\alpha_k^{p+1}L_f}{p+1}\|y_k - x_k\|^{p+1}. \end{aligned} \quad (4.4)$$

and

$$\begin{aligned} g_{\beta_k}(Ax_{k+1}) &\leq g_{\beta_k}(Ax_k) + \langle \nabla(g_{\beta_k} \circ A)(x_k), x_{k+1} - x_k \rangle + \frac{1}{2\beta_k}\|A\|^2\|x_{k+1} - x_k\|^2 \\ &= g_{\beta_k}(Ax_k) + \alpha_k \langle \nabla(g_{\beta_k} \circ A)(x_k), y_k - x_k \rangle + \frac{\alpha_k^2}{2\beta_k}\|A\|^2\|y_k - x_k\|^2. \end{aligned} \quad (4.5)$$

Applying Lemma 4.2 to  $g_{\beta_k}(Ax_{k+1})$  yields

$$\begin{aligned} g_{\beta_{k+1}}(Ax_{k+1}) &\leq g_{\beta_k}(Ax_k) + \alpha_k \langle \nabla(g_{\beta_k} \circ A)(x_k), y_k - x_k \rangle + \frac{\alpha_k^2}{2\beta_k}\|A\|^2\|y_k - x_k\|^2 \\ &\quad + \frac{1}{2}(\beta_k - \beta_{k+1})L_g^2. \end{aligned}$$

Adding (4.4) and (4.5) we obtain that

$$\begin{aligned} F_{k+1}(x_{k+1}) &\leq F_k(x_k) + \alpha_k \langle \nabla F_k(x_k), y_k - x_k \rangle + \frac{\alpha_k^{p+1}L_f}{p+1}\|y_k - x_k\|^{p+1} \\ &\quad + \frac{\alpha_k^2}{2\beta_k}\|A\|^2\|y_k - x_k\|^2 + \frac{1}{2}(\beta_k - \beta_{k+1})L_g^2. \end{aligned} \quad (4.6)$$

By Theorem 2.3 and Proposition 2.2,  $y_k \in \text{dom } \omega$  for every  $k \in \mathbb{N}$ . Since  $\alpha_k \in [0, 1]$ , applying an induction argument yields that  $x_k \in \text{dom } \omega$ , for every  $k \in \mathbb{N}$ . Thus, the values of the objective generated by the algorithm are finite. From the construction of  $y_k$  in steps (I), (II) and Theorem 2.3, we obtain that, for every  $x \in \text{dom } \omega$ ,

$$\langle y_k, -\nabla F_k(x_k) \rangle - \omega(y_k) \geq \langle x, -\nabla F_k(x_k) \rangle - \omega(x)$$

and hence that

$$\begin{aligned} \langle y_k - x_k, -\nabla F_k(x_k) \rangle - \omega(y_k) &\geq \langle x - x_k, -\nabla F_k(x_k) \rangle - \omega(x) \\ &\geq F_k(x_k) - F_k(x) - \omega(x). \end{aligned}$$

Applying Lemma 4.2 to  $g_{\beta_k}(Ax)$  yields

$$\langle y_k - x_k, -\nabla F_k(x_k) \rangle - \omega(y_k) \geq F_k(x_k) - F(x) - \omega(x)$$

and, therefore,

$$\alpha_k \langle y_k - x_k, -\nabla F_k(x_k) \rangle - \alpha_k \omega(y_k) \geq \alpha_k F_k(x_k) - \alpha_k (F(x) + \omega(x)). \quad (4.7)$$

Adding (4.6) and (4.7), we obtain that

$$\begin{aligned} F_{k+1}(x_{k+1}) + \alpha_k F_k(x_k) - \alpha_k (F(x) + \omega(x)) &\leq F_k(x_k) - \alpha_k \omega(y_k) \\ &+ \frac{\alpha_k^{p+1} L_f}{p+1} \|y_k - x_k\|^{p+1} + \frac{\alpha_k^2}{2\beta_k} \|A\|^2 \|y_k - x_k\|^2 + \frac{1}{2} (\beta_k - \beta_{k+1}) L_g^2 \end{aligned}$$

or that

$$\begin{aligned} F_{k+1}(x_{k+1}) + \omega(x_{k+1}) - F(x) - \omega(x) &\leq (1 - \alpha_k) (F_k(x_k) - F(x) - \omega(x)) + \omega(x_{k+1}) \\ &- \alpha_k \omega(y_k) + \frac{\alpha_k^{p+1} L_f}{p+1} \|y_k - x_k\|^{p+1} + \frac{\alpha_k^2}{2\beta_k} \|A\|^2 \|y_k - x_k\|^2 + \frac{1}{2} (\beta_k - \beta_{k+1}) L_g^2 \\ &\leq (1 - \alpha_k) (F_k(x_k) + \omega(x_k) - F(x) - \omega(x)) + \frac{\alpha_k^{p+1} L_f}{p+1} \|y_k - x_k\|^{p+1} \\ &\quad + \frac{\alpha_k^2}{2\beta_k} \|A\|^2 \|y_k - x_k\|^2 + \frac{1}{2} (\beta_k - \beta_{k+1}) L_g^2, \end{aligned}$$

where the last step uses the convexity of  $\omega$  and (III). Since  $x_k, y_k \in \text{dom } \omega$ , for every  $k \in \mathbb{N}$ , it follows that  $\|y_k\|, \|x_k\| \leq \rho$  and hence that (4.3) holds.  $\square$

**Corollary 4.4.** *Suppose that,  $\alpha_1 = 1$ ,  $\alpha_k \in [0, 1]$  and  $\beta_k \geq \beta_{k+1} > 0$ , for every  $k \in \mathbb{N}$ . Let  $P_j = \prod_{i=j+1}^k (1 - \alpha_i)$ , for every  $j \in \{1, \dots, k\}$ . Then*

$$\begin{aligned} f(x_{k+1}) + g(Ax_{k+1}) + \omega(x_{k+1}) - f(x) - g(Ax) - \omega(x) &\leq \frac{(2\rho)^{p+1} L_f}{p+1} \sum_{j=1}^k P_j \alpha_j^{p+1} \\ &+ 2\|A\|^2 \rho^2 \sum_{j=1}^k P_j \frac{\alpha_j^2}{\beta_j} + \frac{1}{2} L_g^2 \sum_{j=1}^k P_j (\beta_j - \beta_{j+1}) + \frac{1}{2} \beta_{k+1} L_g^2 \end{aligned}$$

for every  $x \in \text{dom } \omega$ ,  $k \in \mathbb{N}$ .

**Proof.** Let  $D_k = f(x_k) + g_{\beta_k}(Ax_k) + \omega(x_k) - f(x) - g(Ax) - \omega(x)$ . Applying Theorem 4.3, we obtain

$$D_{j+1} \leq (1 - \alpha_j) D_j + \frac{(2\rho)^{p+1} L_f}{p+1} \alpha_j^{p+1} + 2\|A\|^2 \rho^2 \frac{\alpha_j^2}{\beta_j} + \frac{1}{2} (\beta_j - \beta_{j+1}) L_g^2$$

for every  $j \in \{1, \dots, k\}$ . Multiplying by  $P_j$  and adding up, we obtain

$$\begin{aligned} D_{k+1} &\leq (1 - \alpha_1)P_1D_1 + \frac{(2\rho)^{p+1}L_f}{p+1} \sum_{j=1}^k P_j \alpha_j^{p+1} + 2\|A\|^2 \rho^2 \sum_{j=1}^k P_j \frac{\alpha_j^2}{\beta_j} \\ &\quad + \frac{1}{2}L_g^2 \sum_{j=1}^k P_j(\beta_j - \beta_{j+1}) \\ &= \frac{(2\rho)^{p+1}L_f}{p+1} \sum_{j=1}^k P_j \alpha_j^{p+1} + 2\|A\|^2 \rho^2 \sum_{j=1}^k P_j \frac{\alpha_j^2}{\beta_j} + \frac{1}{2}L_g^2 \sum_{j=1}^k P_j(\beta_j - \beta_{j+1}) \end{aligned}$$

Applying Lemma 4.2 to  $g_{\beta_{k+1}}(Ax_{k+1})$  the assertion follows.  $\square$

**Corollary 4.5.** *If  $\alpha_k = \frac{2}{k+1}$ ,  $\beta > 0$  and  $\beta_k = \frac{\beta}{\sqrt{k}}$ , for every  $k \in \mathbb{N}$ , then*

$$\begin{aligned} f(x_{k+1}) + g(Ax_{k+1}) + \omega(x_{k+1}) - f(x) - g(Ax) - \omega(x) &\leq \\ &\frac{(4\rho)^{p+1}L_f}{(p+1)(k+1)^p} + \frac{8\rho^2\|A\|^2}{\beta\sqrt{k+1}} + \frac{1}{2}L_g^2\beta \frac{\sqrt{k+2}}{k} + \frac{L_g^2\beta}{2\sqrt{k+1}} \end{aligned}$$

for every  $x \in \text{dom } \omega$ ,  $k \in \mathbb{N}$ .

**Proof.** It follows easily from Corollary 4.4 and the computation  $P_j = \frac{j(j+1)}{k(k+1)}$ .  $\square$

We notice that when  $p \geq \frac{1}{2}$  the asymptotic rate does not depend on  $p$  and translates to  $\mathcal{O}\left(\frac{1}{\varepsilon^2}\right)$  iterations, if  $\varepsilon$  is the precision in terms of the objective function. This rate of convergence is an order of magnitude slower than the rate for the standard conditional gradient algorithm (Corollary 3.2). Thus, the extended flexibility of handling multiple additional penalties (function  $g$ ) and the Moreau smoothing incur a cost in terms of iterations. In other words, the class of optimization problems to which the hybrid algorithm applies is significantly larger than that of the standard algorithm 1 and a deterioration in the rate of convergence is inevitable. When  $0 < p < \frac{1}{2}$ , the bound is dominated by the term involving  $p$  and the number of iterations required grows as  $\mathcal{O}\left(\varepsilon^{-\frac{1}{p}}\right)$ .

If there is no  $g \circ A$  term ( $A = 0, g = 0$ ) then the algorithm becomes the standard conditional gradient and the corollary reduces to known bounds for standard conditional gradient methods. The number of iterations grows as  $\mathcal{O}\left(\varepsilon^{-\frac{1}{p}}\right)$ , which ranges from  $\mathcal{O}\left(\frac{1}{\varepsilon}\right)$  (for  $p = 1$ ) to impractical when  $f$  is too “close” to a Lipschitz continuous function ( $p \simeq 0$ ).

The rate in Corollary 4.5 is also slower than the  $\mathcal{O}\left(\frac{1}{\varepsilon}\right)$  rates obtained with smoothing methods, such as [44, 42, 25, 7]. However, smoothing methods require a more powerful computational oracle (the proximity operator of  $\omega$  instead of the dual subgradient) and hence may be inapplicable in problems like those involving very large matrices, because computation of  $\text{prox}_\omega$  may not scale well. Another

$\mathcal{O}(\frac{1}{\varepsilon^2})$  alternative is subgradient methods, but these may be inapplicable too for similar reasons. For example, the subgradient of the trace norm as either a penalty term or a constraint requires a full singular value decomposition.

In addition, like other conditional gradient methods or greedy methods and matching pursuits, the HCGS algorithm 2 builds a parsimonious solution in additive fashion rather than starting from a complex solution and then simplifying it. This feature may be desirable in itself whenever a parsimonious solution is sought. For example, in many cases it is more important to obtain a sparse or low rank estimate of the solution rather than a more accurate one with many small nonzero components or singular values. In machine learning problems, especially, this is frequently the case since the optimization objective is just an approximation of the ideal measure of expected risk [51]. Another advantage of such algorithmic schemes is computational. In sparse estimation problems regularized with an  $\ell_1$  constraint, the data matrix or the dictionary may be huge and hence computation of  $\nabla f(x)$  may be feasible only for sparse vectors  $x$  (when  $f$  is a quadratic function). Moreover, such a computation can be done efficiently since the gradient from the previous iteration can be reused, due to update (III).

### 4.3 Minimization of Lipschitz Continuous Functions

A special case of particular interest occurs when  $f = 0$ , that is, when there is no smooth part. Then HCGS solves the optimization problem

$$\min \{g(Ax) + \omega(x) : x \in \mathcal{H}\} \quad (4.8)$$

under Assumption 4.1 as before. Namely, the objective function consists of a Lipschitz term  $g$  and a generic term  $\omega$  defined on a bounded domain. For example, such a problem is the minimization of a Lipschitz continuous function over a bounded domain. More generally,  $g \circ A$  may incorporate a sum of multiple Lipschitz continuous penalties.

The HCGS algorithm specified to problem (4.8) is the same as Algorithm 2 with  $\nabla f(x_k)$  removed. In this way, the computational model of conditional gradient methods extends from minimization of smooth functions to minimization of Lipschitz continuous functions. Moreover, the convergence rate deteriorates from  $\mathcal{O}(\frac{1}{\varepsilon})$  for smooth functions to  $\mathcal{O}(\frac{1}{\varepsilon^2})$  for Lipschitz functions – which is not surprising, since the latter are in general more difficult to optimize than the former. This fact has been shown recently by Lan for several conditional gradient algorithms [34]. Lan has also shown that these rates coincide with the lower complexity bounds for a family of algorithms involving  $\partial\omega^*$  oracles. The above fact is also intriguing in view of the analogy to the results known about Nesterov’s proximal methods [40]. Those methods, under a more powerful computational oracle for  $\omega$ , exhibit an  $\mathcal{O}(\frac{1}{\sqrt{\varepsilon}})$  rate when  $g$  is smooth versus an  $\mathcal{O}(\frac{1}{\varepsilon})$  rate when  $g$  is Lipschitz continuous.

## 4.4 Implementation Details

It is worth noting that the HCGS algorithm does not require knowledge of the Lipschitz constants  $L_f, L_g$  and can be implemented with an arbitrary choice of  $\beta$ . An alternative is to optimize the bound in Corollary 4.5 with respect to  $\beta$ , which gives an optimal choice of  $\frac{2\sqrt{2}\rho\|A\|}{L_g}$ , asymptotically. If the desired accuracy  $\varepsilon$  can be specified in advance, then the optimal  $\beta$  will also depend on  $\varepsilon$ . Computing such a  $\beta$  value is possible only if the Lipschitz constant and bound of the optimization problem are available, but for regularization problems these constants can be computed from the regularization parameters.

For  $p \geq \frac{1}{2}$ , these two constants,  $\rho$  and  $L_g$ , have the largest influence in the convergence rate, since they appear in the  $\mathcal{O}\left(\frac{1}{\sqrt{k}}\right)$  terms that dominate the bound. The constant  $\rho$  cannot be changed, since it is a property of the feasibility domain. However  $L_g$  can be reduced by rescaling the objective function and hence it can become independent of the dimensionality of the problem.

Some care may be needed to tackle numerical issues arising from very small values of  $\beta_k$  as  $k$  becomes large. These issues affect only step (I), whereas the computation of  $y_k$  in step (II) remains always inside the  $\rho$ -ball, since  $\omega^*$  is  $\rho$ -Lipschitz continuous. Moreover, for large  $k$ , the past estimates dominate the update (III) and hence the effect of any numerical issues diminishes as  $k$  grows.

## 5 Applications

We now instantiate the HCGS algorithm 2 to some special cases which appear in applications and we present the corresponding algorithms. These examples are only a sample and do not cover the whole range of possible applications. First, consider the problem of learning a *sparse and low rank matrix* by regularization with the  $\ell_1$  norm and a trace norm constraint [46],

$$\min\{f(X) + \lambda\|X\|_1 : \|X\|_{tr} \leq B, X \in \mathbb{R}^{n \times n}\}, \quad (5.1)$$

where  $\|\cdot\|_1$  denotes the elementwise  $\ell_1$  norm of a matrix and  $\|\cdot\|_{tr}$  the trace norm (or nuclear norm). The strongly smooth function  $f$  expresses an error term (where the dependence on the data is absorbed in  $f$ ) and may arise by using, for example, the square loss or the logistic loss. This setting has been proposed for applications such as graph denoising or prediction of links on a social network. The resulting algorithm (Algorithm 3) depends on the proximity operator of the  $\ell_1$  norm, also known as the soft thresholding operator,

$$\mathcal{S}(X; \gamma) = \text{sgn}(X) \odot (|X| - \gamma)_+,$$

where  $\text{sgn}, \odot, |\cdot|$  denote elementwise sign, multiplication and absolute value on matrices and  $(\cdot)_+$  the positive part elementwise.

Note that the same algorithm can be used for solving a variation of (5.1) that restricts the optimization to the space of symmetric matrices. This may



---

**Algorithm 3** Hybrid algorithm for sparse - low rank problems.

---

**Input**  $X_1 \in \mathbb{R}^{n \times n}$  such that  $\|X_1\|_{tr} \leq B$   
**for**  $k = 1, 2, \dots$  **do**  
 $Z_k \leftarrow -\nabla f(X_k) - \frac{1}{\beta_k} X_k + \frac{1}{\beta_k} \mathcal{S}(X_k; \beta_k \lambda)$   
 $(u_k, v_k) \leftarrow$  a left and right pair of singular vectors of  $Z_k$  corresponding to the largest singular value  
 $Y_k \leftarrow B u_k v_k^\top$   
 $X_{k+1} \leftarrow (1 - \alpha_k) X_k + \alpha_k Y_k$   
**end for**

---



---

**Algorithm 4** Hybrid algorithm for sparse PCA relaxation.

---

**Input**  $X_1 \in \mathbb{R}^{n \times n}$  such that  $\text{tr}(X_1) = 1, X_1 \succeq 0$   
**for**  $k = 1, 2, \dots$  **do**  
 $Z_k \leftarrow C - \frac{1}{\beta_k} X_k + \frac{1}{\beta_k} \mathcal{S}(X_k; \beta_k \lambda)$   
 $u_k \leftarrow$  a dominant eigenvector of  $Z_k$   
 $Y_k \leftarrow u_k u_k^\top$   
 $X_{k+1} \leftarrow (1 - \alpha_k) X_k + \alpha_k Y_k$   
**end for**

---

occur, for example, when learning the adjacency matrix of an undirected graph. One should ensure, however, that the initial matrix  $X_1$  is symmetric.

A problem which shares some similarities with the previous one is the convex relaxation of *sparse PCA* proposed in [2],

$$\max\{\langle C, X \rangle - \lambda \|X\|_1 : \text{tr}(X) = 1, X \succeq 0, X \in \mathbb{R}^{n \times n}\}. \quad (5.2)$$

Solving this optimization problem can be used for finding a dominant sparse eigenvector of  $C$ , which is a prescribed  $n \times n$  symmetric matrix. The problem falls under the framework (4.1) with  $f$  being a linear function and  $\omega$  the indicator function of the (bounded) spectrahedron  $\{X \in \mathbb{R}^{n \times n} : \text{tr}(X) = 1, X \succeq 0\}$ . Computation of a dual subgradient amounts to computing a solution of the problem

$$\max\{\langle Y, Z \rangle : \text{tr}(Y) = 1, Y \succeq 0, Y \in \mathbb{R}^{n \times n}\}$$

for a given symmetric matrix  $Z \in \mathbb{R}^{n \times n}$ . It is easy to see that this computation requires a dominant eigenvector of  $Z$ . This results in Algorithm 4.

A related problem is to restrict the sparse - low rank optimization (5.1) to the cone of positive semidefinite matrices. This problem has been proposed for estimating a covariance matrix in [46]. Since the trace norm is equal to the trace on the positive semidefinite cone, the algorithm is similar to Algorithm 4. The only differences are the initialization, a general smooth function  $f$  and a factor of  $B$  in the update of  $Y_k$ .

A third example of an optimization problem that falls under our framework

---

**Algorithm 5** Hybrid algorithm for sparse multicomposite problems (5.3).

---

**Input**  $x_1 = Be_i$  for some  $i \in \{1, \dots, d\}$   
**for**  $k = 1, 2, \dots$  **do**  
 $z_k \leftarrow -Qx_k - c - \frac{1}{\beta_k}A^*Ax_k + \frac{1}{\beta_k}A^*\text{prox}_{\beta_k g}(Ax_k)$   
 $y_k \leftarrow B \text{sgn}((z_k)_j)e_j$ , where  $j \in \text{argmax}_{i=1}^d |(z_k)_i|$   
 $x_{k+1} \leftarrow (1 - \alpha_k)x_k + \alpha_k y_k$   
**end for**

---

is a regularization problem with  $\ell_1$  and additional penalties,

$$\min \left\{ \frac{1}{2} \langle x, Qx \rangle + \langle c, x \rangle + g(Ax) : \|x\|_1 \leq B, x \in \mathbb{R}^d \right\}. \quad (5.3)$$

Here  $Q \in \mathbb{R}^{d \times d}$  is a prescribed positive semidefinite matrix,  $c \in \mathbb{R}^d$  a prescribed vector and  $g, A$  satisfy Assumption 4.1. For example, (5.3) could arise from an estimation or learning problem, the quadratic part corresponding to the data fit term. The  $\ell_1$  constraint is used to favor sparse solutions. The penalty terms  $g \circ A$  may involve multiple norms whose proximity operator is simple to compute, such as the group Lasso norm [57], total variation norms [50] etc.

The hybrid method, specialized to such problems, is shown in Algorithm 5. In general, several other algorithms may be used for solving problems like (5.3) (smoothing, Douglas-Rachford, subgradient methods etc.), but here we are interested in cases with *very large dimensionality*  $d$ . In such cases, computation of the gradient at an arbitrary vector is  $\mathcal{O}(d^2)$  and very costly. On the other side, in the HCGS algorithm,  $x_{k+1}$  is  $(k+1)$ -sparse and computing the new gradient  $Qx_{k+1}$  can be done efficiently by keeping  $Qx_k$  in memory and computing  $Qy_k$ , which is proportional to the  $j$ -th column of  $Q$ . The latter requires only  $\mathcal{O}(d)$  operations, or  $\mathcal{O}(dm)$  if  $Q$  is the square of an  $m \times d$  data matrix. Thus, HCGS can be applied to such problems at a smaller cost, by starting with an initial cardinality-one vector and stopping before  $k$  becomes too large.

There is also an interesting interpretation of Algorithm 5 as an extension of *matching pursuits* [36, 55] to problems with multiple penalties. Assuming that  $Q$  is the square of a matrix of dictionary elements (or more generally a Gram matrix of elements from a Hilbert space), then the algorithm shares similarities with orthogonal matching pursuit (OMP). Indeed, such a connection has already been observed for the standard conjugate gradient (which corresponds to absence of the  $g \circ A$  term) [28, 27], the main difference from OMP being in the update of  $x_{k+1}$ . Similarly, the HCGS algorithm 5 could be phrased as an extension of OMP that imposes additional penalties  $g \circ A$ , besides sparsity, on the coefficients of the atoms. For example,  $g \circ A$  could involve *structured sparsity* penalties (such as penalties for group, hierarchical or graph sparsity) and then HCGS would yield a scalable alternative to structured variants of OMP [26] or proximal methods for structured sparsity [35].

## 6 Simulations

### 6.1 Simultaneous Sparse and Low Rank Regularization

In this section we focus on testing Algorithm 3 (HCGS) on the estimation of simultaneously sparse and low rank matrices.<sup>4</sup> Our aim is to compare the procedure with the proximal algorithms proposed, for the same task, in [46]. The experiments illustrate the fact that HCGS scales better than the SVD-based alternatives.

We considered the task of recovering a matrix from a subset of its entries. To this end in each simulation we generated two  $N \times 5$  random matrices with entries drawn from the uniform distribution. 90% of the entries corresponding to a subset of uniformly distributed indices were then set to zero. The resulting matrices, denoted by  $U$  and  $V$ , were then used to obtain a sparse and low rank matrix  $UV^\top$ . This matrix was corrupted by zero-mean Gaussian noise with variance  $\sigma^2 = 10^{-4}$  to obtain the observation matrix  $Y$ . A fraction  $f \in \{0.05, 0.4\}$  of entries of  $Y$  were used for recovery; see Figure 1 for an illustration.

We compared HCGS with the two algorithms proposed in [46], namely *generalized forward-backward* (GFB) [45] and *incremental proximal descent* (IPD) [8]. Both these algorithms solve a convex matrix recovery problem that aims at finding a matrix that is simultaneously low rank and sparse. This problem is:

$$\min \left\{ J(X) := \frac{1}{2p} \|\Omega(Y - X)\|_F^2 + \lambda_1 \|X\|_1 + \lambda_2 \|X\|_{tr} : X \in \mathbb{R}^{N \times N} \right\} \quad (6.1)$$

where  $\|\cdot\|_F$  is the Frobenius norm,  $p$  is the number of observed entries and  $\Omega : \mathbb{R}^{N \times N} \rightarrow \mathbb{R}^{N \times N}$  is the sampling operator defined entry-wise by  $\Omega(X)_{ij} = X_{ij}$  if the entry indexed by  $(i, j)$  is observed,  $\Omega(X)_{ij} = 0$  otherwise. In our experiments we set  $\lambda_1 = \frac{1}{N^2}$ ,  $\lambda_2 = \frac{10^{-3}}{N^2}$  and used GFB and IPD to obtain optimal estimates  $\hat{X}_{\text{GFB}}$  and  $\hat{X}_{\text{IPD}}$ , respectively. We then set  $\tau := \|\hat{X}_{\text{GFB}}\|_{tr}$  and used HCGS to solve the constrained formulation equivalent to (6.1), namely:


$$\min \left\{ \frac{1}{2p} \|\Omega(Y - X)\|_F^2 + \lambda_1 \|X\|_1 : \|X\|_{tr} \leq \tau, X \in \mathbb{R}^{N \times N} \right\}. \quad (6.2)$$

The comparisons were performed on an Intel Xeon with 8 cores and 24GB of memory.


Figure 2 shows the evolution of objective values for  $N = 400$ . Note that for the sake of comparison we have reported the objective of the optimization problem (6.1) even though HCGS actually solves the equivalent problem in (6.2). The same applies to the attained objective function value  $J_{k^*}$  in Table 1. In this table we have also compared the different algorithms in terms of CPU time<sup>5</sup>, relative change in the objective function and number of iterations upon

<sup>4</sup>Code is available at <http://cvm.eep.fr/personnel/andreas/code/index.html>


<sup>5</sup>In Figure 2b the CPU times include also the evaluation of the objective value in (6.1) which requires computing the singular values of the estimate. In the case of HCGS, this is required only for the sake of comparison with GFB and IPD. In contrast, for Table 1 the objective value of (6.1) is computed only upon termination.




(a)



(b)



(c)



(d)

Figure 1: An illustration of the synthetic problem for  $N = 200$ . The generating matrix, simultaneously sparse and low-rank (a), a mask with the observed entries, in white; 40% of the total number of entries are observed (b), the matrix estimated by HCGS (c), the leading singular values for the generating/observation/estimated matrix (d).




Figure 2: Comparison of objective values for  $N = 400$ , (a) as a function of the iteration count, and (b) as a function of time.




Figure 3: Required time for convergence as a function of  $N$  (40% observed entries), for the sparse - low rank experiment.

termination. In all the cases we terminated the algorithms at iteration  $k^*$  when the relative change  $r_{k^*}$  in the objective value:

$$r_{k^*} = \left| \frac{f_{k^*} - f_{k^*-1}}{f_{k^*-1}} \right| \quad (6.3)$$

was less than  $10^{-7}$ . Note that  $f$  in (6.3) refers to the objective function actually minimized by each algorithm; this is not necessarily the objective function  $J$  in (6.1). Figure 3 shows the time complexity as a function of  $N$ . Finally, in Table 2 we have reported the average time per iteration as a function of  $N$ .

From these figures and tables, we see that the running time of HCGS scales as  $\mathcal{O}(N^2)$  with the matrix size, whereas both GFB and IPD scale as  $\mathcal{O}(N^3)$ .

## 6.2 Sparse PCA

The second set of simulations assesses the computational efficiency of HCGS on the convex relaxation of sparse PCA (5.2). Similar to [2], we generated random matrices  $C$  as follows. For each size  $n$ , we drew an  $n \times n$  matrix  $U$  with uniformly distributed elements in  $[0, 1]$ . Then we generated a vector  $v \in \mathbb{R}^n$  from the uniform distribution and set a random 90% of its components to zero. We then set

$$C = UU^\top + 10vv^\top.$$

We solved (5.2) for  $\lambda = 1$  with HCGS (Algorithm 4) and the Nesterov smoothing method of [2] which optimizes the dual problem of (5.2). We implemented both algorithms in Matlab and used a cluster with 24 cores and sufficient memory. For HCGS we used the power method with a tolerance of  $10^{-6}$ , for computing dominant eigenvectors. We also rescaled the objective function by  $n$  in order to keep  $L_g$  small enough (see Section 4.4). For Nesterov smoothing we used  $\mu = \frac{10^{-6}}{2 \log n}$ .

In Figure 6.2, we plot the computational times required to attain relative change of  $10^{-5}$  in the objective. We note that the objective functions are different, since HCGS optimizes (5.2) whereas the method of [2] optimizes the dual problem. In fact, we have verified that the duality gap estimates are consistently larger for the latter and hence the running times for Nesterov smoothing are optimistic. We observe that the running time scales roughly as  $\mathcal{O}(n^2)$  for HCGS whereas Nesterov smoothing scales worse than  $\mathcal{O}(n^3)$ .

## 7 Conclusion

We have studied the hybrid conditional gradient - smoothing algorithm (HCGS) for solving composite convex optimization problems which contain several terms over a bounded set. Examples of these include regularization problems with several norms as penalties and a norm constraint. HCGS extends conditional gradient methods to cases with multiple nonsmooth terms, in which standard conditional gradient methods may be difficult to apply. The HCGS algorithm

Table 1: Comparison of different algorithms for convex matrix recovery.

5% observed entries

N		$J_{k^*} (\times 10^{-4})$	$r_{k^*} (\times 10^{-8})$	time (s)	$k^*$
50	HCGS	6.77	8.35	2.38	1605
	GFB	6.87	5.15	0.09	45
	IPD	6.76	9.07	0.09	50
100	HCGS	4.41	7.52	3.60	1462
	GFB	4.45	9.67	2.57	385
	IPD	4.40	9.80	2.5	388
200	HCGS	4.16	2.42	66	3308
	GFB	4.17	9.95	113	4136
	IPD	4.16	9.83	123	4645
400	HCGS	2.98	3.35	176	4555
	GFB	2.98	9.99	2333	15241
	IPD	2.98	9.99	2389	15665
600	HCGS	2.46	7.15	158	2797
	GFB	2.45	9.99	13157	36049
	IPD	2.45	9.99	13080	36408
800	HCGS	2.20	5.47	478	5197
	GFB	2.20	9.99	40779	61299
	IPD	2.20	9.99	41263	61529

40% observed entries

N		$J_{k^*} (\times 10^{-4})$	$r_{k^*} (\times 10^{-8})$	time (s)	$k^*$
50	HCGS	18.51	8.85	0.62	427
	GFB	18.66	8.75	1.82	751
	IPD	18.52	9.96	2.17	1004
100	HCGS	12.09	4.96	2.25	835
	GFB	12.15	9.84	12.1	1681
	IPD	12.10	6.50	11.8	1697
200	HCGS	7.65	9.64	21	1410
	GFB	7.66	9.98	134	3521
	IPD	7.65	7.66	112	3033
400	HCGS	4.39	2.54	40	1281
	GFB	4.39	9.96	1559	7379
	IPD	4.39	9.46	1580	7515
600	HCGS	3.41	7.65	43	760
	GFB	3.41	9.98	5664	10793
	IPD	3.41	9.99	5446	10841
800	HCGS	2.68	1.37	148	1378
	GFB	2.68	9.98	14897	15615
	IPD	2.68	9.99	11242	15647
1600	HCGS	1.62	9.91	929	2931
	GFB	1.62	9.99	119724	36573
	IPD	1.62	9.99	118223	36583

Table 2: Average time (in seconds) per iteration as a function of  $N$  (40% observed entries), for the sparse - low rank experiment.

	$N$				
	200	400	600	800	1600
HCGS	0.0155	0.032	0.058	0.107	0.317
GFB	0.038	0.211	0.524	0.954	3.273
IPD	0.037	0.210	0.502	0.718	3.231




Figure 4: Computational time (in seconds) versus matrix size for the sparse PCA experiment.

borrow techniques from smoothing proximal methods and requires first-order computations (subgradients and proximity operations). Moreover, it exhibits convergence in terms of the objective values at an  $\mathcal{O}(\frac{1}{\varepsilon^2})$  rate of iterations. Unlike proximal methods, HCGS benefits from the advantages of conditional gradient methods, which render it more efficient on certain large scale optimization problems. We have demonstrated these advantages with simulations on two matrix optimization problems: regularization of matrices with combined  $\ell_1$  and trace norm penalties; and a convex relaxation of sparse PCA.

## Acknowledgments

The research leading to the above results has received funding from the European Union Seventh Framework Programme (FP7 2007-2013) under grant agreement No. 246556. The research has also received funding from the European Union Seventh Framework Programme (FP7 2007-2013) under grant agreement No. 246556. Research supported also by: Research Council KUL: GOA/10/09



MaNet, PFV/10/002 (OPTEC), several PhD/postdoc and fellow grants; Flemish Government:IOF: IOF/KP/SCORES4CHEM, FWO: PhD/postdoc grants, projects: G.0377.12 (Structured systems), G.083014N (Block term decompositions), G.088114N (Tensor based data similarity), IWT: PhD Grants, projects: SBO POM, EUROSTARS SMART, iMinds 2013, Belgian Federal Science Policy Oce: IUAP P7/19 (DYSCO, Dynamical systems, control and optimization, 2012-2017), EU: FP7-SADCO (MC ITN- 264735), ERC AdG A-DATADRIVE-B (290923), COST: Action ICO806: IntelliCIS.

## References

- [1] J. Abernethy, F. Bach, T. Evgeniou, and J-P. Vert. A new approach to collaborative filtering: Operator estimation with spectral regularization. *Journal of Machine Learning Research*, 10:803–826, 2009.
- [2] A. d’ Aspremont, L. El Ghaoui, M. I. Jordan, and G. R. G. Lanckriet. A direct formulation for sparse PCA using semidefinite programming. *SIAM Review*, pages 434–448, 2007.
- [3] F. Bach. Duality between subgradient and conditional gradient methods. *arXiv preprint arXiv:1211.6302*, 2012.
- [4] H. H. Bauschke and P. L. Combettes. *Convex Analysis and Monotone Operator Theory in Hilbert Spaces*. CMS Books in Mathematics. Springer, 2011.
- [5] A. Beck and M. Teboulle. Mirror descent and nonlinear projected subgradient methods for convex optimization. *Operations Research Letters*, 31(3):167–175, 2003.
- [6] A. Beck and M. Teboulle. A fast iterative shrinkage-thresholding algorithm for linear inverse problems. *SIAM Journal of Imaging Sciences*, 2(1):183–202, 2009.
- [7] A. Beck and M. Teboulle. Smoothing and first order methods: a unified framework. *SIAM Journal on Optimization*, 22(2):557–580, 2012.
- [8] D. P. Bertsekas. Incremental gradient, subgradient, and proximal methods for convex optimization: A survey. In S. Sra, S. Nowozin, and S. J. Wright, editors, *Optimization for Machine Learning*, pages 85–119. MIT Press, 2011.
- [9] M. D. Canon and C. D. Cullum. A tight upper bound on the rate of convergence of Frank-Wolfe algorithm. *SIAM Journal on Control*, 6(4):509–516, 1968.
- [10] K. L. Clarkson. Coresets, sparse greedy approximation, and the Frank-Wolfe algorithm. *ACM Transactions on Algorithms*, 6(4):63:1–63:30, 2010.

- [11] P. L. Combettes and J. C. Pesquet. Proximal splitting methods in signal processing. In *Fixed-Point Algorithms for Inverse Problems in Science and Engineering*, pages 185–212. Springer, 2011.
- [12] P. L. Combettes and V. R. Wajs. Signal recovery by proximal forward-backward splitting. *Multiscale Modeling and Simulation*, 4(4):1168–1200, 2006.
- [13] V. F. Dem’yanov and A. B. Pevnyi. Some estimates in minimax problems. *Cybernetics and Systems Analysis*, 8(1):116–123, 1972.
- [14] V. F. Dem’yanov and A. M. Rubinov. *Approximate methods in optimization problems*, volume 32. Elsevier, 1970.
- [15] J. C. Dunn. Rates of convergence for conditional gradient algorithms near singular and nonsingular extremals. *SIAM Journal on Control and Optimization*, 17(2):187–211, 1979.
- [16] J. C. Dunn. Convergence rates for conditional gradient sequences generated by implicit step length rules. *SIAM Journal on Control and Optimization*, 18(5):473–487, 1980.
- [17] J. Eckstein and D. P. Bertsekas. On the Douglas-Rachford splitting method and the proximal point algorithm for maximal monotone operators. *Mathematical Programming*, 55(1-3):293–318, 1992.
- [18] M. Frank and P. Wolfe. An algorithm for quadratic programming. *Naval research logistics quarterly*, 3(1-2):95–110, 1956.
- [19] R. M. Freund and P. Grigas. New analysis and results for the conditional gradient method. Technical report, Massachusetts Institute of Technology, Operations Research Center, 2013.
- [20] D. Garber and E. Hazan. Approximating semidefinite programs in sublinear time. In *Advances in Neural Information Processing Systems 24*, pages 1080–1088, 2011.
- [21] E. G. Gilbert. An iterative procedure for computing the minimum of a quadratic form on a convex set. *SIAM Journal on Control*, 4(1):61–80, 1966.
- [22] Z. Harchaoui, A. Juditsky, and A. Nemirovski. Conditional gradient algorithms for machine learning. In *NIPS Workshop on Optimization for Machine Learning*, 2012.
- [23] E. Hazan. Sparse approximate solutions to semidefinite programs. In *Proceedings of the 8th Latin American conference on Theoretical informatics, LATIN’08*, pages 306–316, 2008.
- [24] E. Hazan and S. Kale. Projection-free online learning. In *Proceedings of the 29th International Conference on Machine Learning*, 2012.

- [25] N. He, A. Juditsky, and A. Nemirovski. Mirror prox algorithm for multi-term composite minimization and alternating directions. *arXiv preprint arXiv:1311.1098*, 2013.
- [26] J. Huang, T. Zhang, and D. Metaxas. Learning with structured sparsity. In *Proceedings of the 26th Annual International Conference on Machine Learning*, pages 417–424. ACM, 2009.
- [27] M. Jaggi. Convex optimization without projection steps. *Arxiv preprint arXiv:1108.1170*, 2011.
- [28] M. Jaggi. Revisiting Frank-Wolfe: Projection-free sparse convex optimization. In *Proceedings of the 30th International Conference on Machine Learning*, pages 427–435, 2013.
- [29] M. Jaggi and M. Sulovský. A simple algorithm for nuclear norm regularized problems. In *Proceedings of the 27th International Conference on Machine Learning*, pages 471–478, 2010.
- [30] A. Juditsky and A. Nemirovski. First order methods for nonsmooth convex large-scale optimization, I: general purpose methods. In *Optimization for Machine Learning*, pages 121–148. MIT Press, 2010.
- [31] H. J. Kelley. Method of gradients. *Academic Press*, 2:1578–1580, 1962.
- [32] V. Kumar. A control averaging technique for solving a class of singular optimal control problems. *International Journal of Control*, 23(3):361–380, 1976.
- [33] S. Lacoste-Julien, M. Jaggi, M. Schmidt, and P. Pletscher. Block-coordinate Frank-Wolfe optimization for structural SVMs. *arXiv preprint arXiv:1207.4747*, 2012.
- [34] G. Lan. The complexity of large-scale convex programming under a linear optimization oracle. *arXiv preprint arXiv:1309.5550*, 2013.
- [35] J. Mairal, R. Jenatton, F. R. Bach, and G. R. Obozinski. Network flow algorithms for structured sparsity. In *Advances in Neural Information Processing Systems*, pages 1558–1566, 2010.
- [36] S. G. Mallat and Z. Zhang. Matching pursuits with time-frequency dictionaries. *IEEE Transactions on Signal Processing*, 41(12):3397–3415, 1993.
- [37] J.J. Moreau. Proximité et dualité dans un espace hilbertien. *Bulletin de la Société Mathématique de France*, 93(2):273–299, 1965.
- [38] A. S. Nemirovsky and D. B. Yudin. *Problem complexity and method efficiency in optimization*. Wiley, 1983.

- [39] Y. Nesterov. A method of solving a convex programming problem with convergence rate  $O(1/k^2)$ . *Soviet Mathematics Doklady*, 27(2):372–376, 1983.
- [40] Y. Nesterov. *Introductory lectures on convex optimization: A basic course*. Springer, 2004.
- [41] Y. Nesterov. Excessive gap technique in nonsmooth convex minimization. *SIAM J. on Optimization*, 16(1):235–249, May 2005.
- [42] Y. Nesterov. Smooth minimization of non-smooth functions. *Mathematical Programming*, 103(1):127–152, 2005.
- [43] Y. Nesterov. Gradient methods for minimizing composite objective function. *CORE Discussion Papers*, 2007.
- [44] F. Orabona, A. Argyriou, and N. Srebro. PRISMA: PROximal Iterative SMOothing Algorithm. arXiv preprint arXiv:1206.2372, 2012.
- [45] H. Raguet, J. Fadili, and G. Peyré. A generalized forward-backward splitting. *SIAM Journal on Imaging Sciences*, 6(3):1199–1226, 2013.
- [46] E. Richard, P.-A. Savalle, and N. Vayatis. Estimation of simultaneously sparse and low rank matrices. In *Proceedings of the 29th International Conference on Machine Learning*, pages 1351–1358, 2012.
- [47] R. T. Rockafellar. Level sets and continuity of conjugate convex functions. *Transactions of the American Mathematical Society*, 123:46–63, 1966.
- [48] R. T. Rockafellar. *Convex Analysis*. Princeton University Press, 1970.
- [49] R. T. Rockafellar. Monotone operators and the proximal point algorithm. *SIAM Journal on Control and Optimization*, 14(5):877–898, 1976.
- [50] L. I. Rudin, S. Osher, and E. Fatemi. Nonlinear total variation based noise removal algorithms. *Physica D: Nonlinear Phenomena*, 60(1):259–268, 1992.
- [51] S. Shalev-Shwartz, N. Srebro, and T. Zhang. Trading accuracy for sparsity in optimization problems with sparsity constraints. *SIAM Journal on Optimization*, 20(6):2807–2832, 2010.
- [52] J. Shawe-Taylor and N. Cristianini. *Kernel Methods for Pattern Analysis*. Cambridge University Press, 2004.
- [53] J. Steinhardt and J. Huggins. A greedy framework for first-order optimization. Preprint.
- [54] A. Tewari, P. K. Ravikumar, and I. S. Dhillon. Greedy algorithms for structurally constrained high dimensional problems. In *Advances in Neural Information Processing Systems 24*, pages 882–890, 2011.

- [55] J. A. Tropp and A. C. Gilbert. Signal recovery from random measurements via orthogonal matching pursuit. *IEEE Transactions on Information Theory*, 53(12):4655–4666, 2007.
- [56] Y. Ying and P. Li. Distance metric learning with eigenvalue optimization. *Journal of Machine Learning Research*, 13:1–26, 2012.
- [57] M. Yuan and Y. Lin. Model selection and estimation in regression with grouped variables. *Journal of the Royal Statistical Society, Series B (Statistical Methodology)*, 68(1):49–67, 2006.